

# EXHIBIT I

**Exhibit 1 – MANTRA I**

'156 Patent

Claim Limitation (Claim 7)	Exemplary Disclosure
[156a] A device comprising:	<p>MANTRA I discloses a device. Specifically, the MANTRA I is a neurocomputer. <i>See, e.g.</i>:</p> <p>“The MANTRA I computer is a massively parallel machine dedicated to neural-network algorithms (Fig. 1). It has been designed to provide the basic operations for the following models: (1) single-layer networks (Perceptron and delta rule); (2) multilayer feedforward networks (back-propagation rule); (3) fully connected recurrent networks (Hopfield model); and (4) self-organizing feature maps (SOFMS; Kohonen model). A description of these algorithms can be found in any classic introductory book on neural networks (e.g., [4 ]). The Kohonen feature maps are used in Section 3 to illustrate how these algorithms are mapped on the system.</p> <p>The MANTRA I accelerator is based on a bidimensional systolic array composed of custom PEs named GENES IV. In the present section, the hardware of the machine is overviewed starting from its system integration in a network of workstations and proceeding down to the internal architecture of the machine and of its computational core.” Thierry Cornu, et al. <i>Design, Implementation, and Test of a Multi-Model Systolic Neural-Network Accelerator</i>, Scientific Programming, Vol. 5, 1996, at 48.</p>

**Exhibit 1 – MANTRA I**

Claim Limitation (Claim 7)	Exemplary Disclosure
	 <p data-bbox="1129 1117 1488 1142"><b>FIGURE 1</b> The MANTRA I system.</p> <p data-bbox="707 1199 1848 1272">Thierry Cornu, et al. <i>Design, Implementation, and Test of a Multi-Model Systolic Neural-Network Accelerator</i>, Scientific Programming, Vol. 5, 1996, at 48.</p> <p data-bbox="707 1313 1890 1419">“The machine, which is currently under test, consists of four printed circuit boards : (1) the processor board (34 chips), (2) the control board (135 chips), (3) the input/output board (465 chips including 14 GACD1 chips), (4) the GENES IV array board (containing a matrix of 10</p>

**Exhibit 1 – MANTRA I**

Claim Limitation (Claim 7)	Exemplary Disclosure
	<p>x10 = 100 GENES IV chips). One of the latter board is required for configurations of the machine up to 20 x 20 = 400 PEs. For larger configurations—up to 40 x 40 = 1600 PEs—four such boards should be interconnected. A more detailed description of the MANTRA I machine can be found in [vir93].” Marc A. Viredaz &amp; Paolo Ienne, <i>MANTRA I: A Systolic Neuro-Computer</i>, In Proceedings of the International Join Conference on Neural Networks, Vol. III, Nagoya, Japan, October 1993, at 3.</p>

## Exhibit 1 – MANTRA I

Claim Limitation (Claim 7)	Exemplary Disclosure
	<p>The diagram illustrates the architecture of the MANTRA I machine. It is organized into two main sections: the SIMD Part and the Exemplary Disclosure section.</p> <p><b>SIMD Part:</b> This section contains several functional blocks:     <ul style="list-style-type: none"> <li><b>GENES Array:</b> Receives inputs <math>U_{OUT}</math> and <math>L_{OUT}</math>, and provides <math>W_{GT_{IN}}</math> to the Weight Input FIFOs.</li> <li><b>Serial to Parallel converters:</b> Convert <math>U_{OUT}</math> to <math>U_{OUL}</math> and <math>L_{OUT}</math> to <math>L_{OUL}</math>.</li> <li><b>Sigma Unit:</b> Receives <math>P</math> and <math>Y</math> from the SIMD Part and provides <math>F_Y</math> to the Delta Unit.</li> <li><b>Conv. Detect. unit:</b> Receives <math>Y</math> and provides <math>OLDX</math> to the SIMD Part.</li> <li><b>X/Y FIFOs and Memories:</b> Store intermediate results <math>X</math> and <math>Y</math>.</li> <li><b>Aux. Y Memory:</b> Stores auxiliary data <math>AY</math>.</li> <li><b>Desired Output FIFO and Memory:</b> Store desired output <math>DES</math>.</li> </ul> </p> <p><b>Exemplary Disclosure:</b> This section shows the overall system architecture:     <ul style="list-style-type: none"> <li><b>SIMD Control:</b> Manages the SIMD part and interacts with the TMS320C40 Microprocessor.</li> <li><b>Instr. FIFO:</b> Provides instructions to the SIMD Control.</li> <li><b>Status Register:</b> Monitors system status.</li> <li><b>Perf. Monitor:</b> Monitors performance metrics.</li> <li><b>Memory:</b> Provides memory access.</li> <li><b>TMS320C40 Microprocessor:</b> Handles control and communication tasks.</li> <li><b>Comm. Links:</b> Connects the SIMD Control to the TMS320C40 Microprocessor.</li> </ul> </p> <p>Communication links are indicated by arrows with widths and data types:     <ul style="list-style-type: none"> <li>From SIMD Control to SIMD Part: <math>WGT_{IN}</math> (40 bits).</li> <li>From SIMD Part to SIMD Control: <math>U_{OUT}</math> (40 bits), <math>L_{OUT}</math> (40 bits), <math>W_{GT}</math> (40 bits), <math>U_{OUL}</math> (40 bits), <math>L_{OUL}</math> (40 bits), <math>OLDX</math> (16 bits), <math>P</math> (40 bits), <math>Y</math> (16 bits), <math>XY</math> (16 bits), <math>AY</math> (16 bits), <math>F_Y</math> (16 bits), <math>Delta</math> (40 bits), <math>Y_1</math> (16 bits), <math>Y_2</math> (16 bits), and <math>DES</math> (16 bits).</li> <li>From SIMD Control to TMS320C40: <math>IIOF_x</math> (32 bits), <math>IIOF_y</math> (32 bits), and <math>D</math> (32 bits).</li> <li>From TMS320C40 to SIMD Control: <math>32</math> (32 bits), <math>32</math> (32 bits), <math>20, 1</math> (20, 1 bits), and <math>32</math> (32 bits).</li> </ul> </p>

Figure 4: Architecture of the MANTRA I machine.

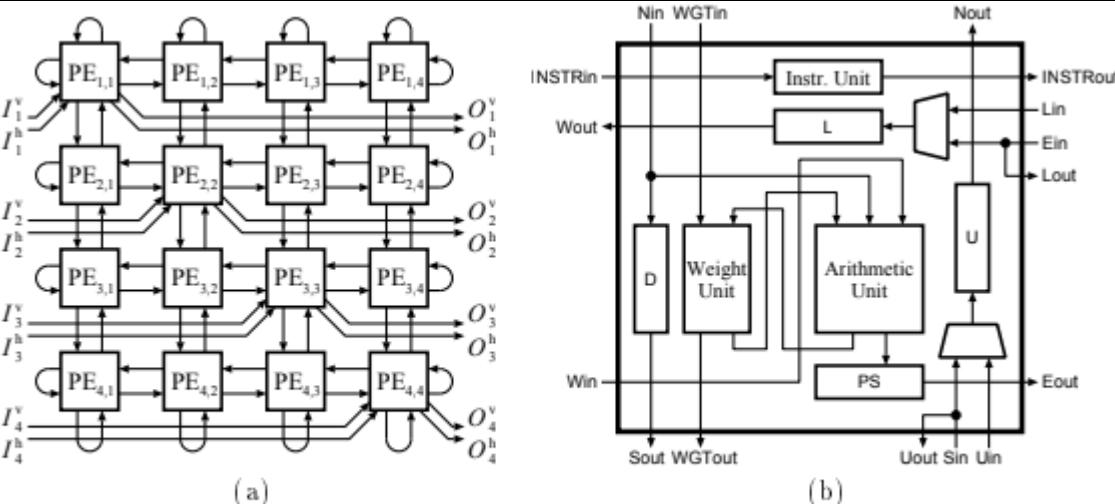
Marc A. Viredaz, *MANTRA I: An SIMD Processor Array for Neural Computation*, Spies P.P. (eds) Europäischer Informatik Kongreß Architektur von Rechensystemen Euro-ARCH, 1993, available at

<https://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.29.9021&rep=rep1&type=pdf>

**Exhibit 1 – MANTRA I**

Claim Limitation (Claim 7)	Exemplary Disclosure
<p>[156b] at least one first low precision high dynamic range (LPHDR) execution unit adapted to execute a first operation on a first input signal representing a first numerical value to produce a first output signal representing a second numerical value,</p>	<p>MANTRA I discloses at least one first low precision high dynamic range (LPHDR) execution unit adapted to execute a first operation on a first input signal representing a first numerical value to produce a first output signal representing a second numerical value. <i>See, e.g.:</i></p> <p>“The systolic array at the heart of the SIMD part of the machine is a square mesh of GENES IV PEs [9], each connected by serial lines to its four neighbours, as shown in Figure 4. All input and output operations are performed by the PEs located on the north-west to south-east diagonal.” Thierry Cornu, et al. <i>Design, Implementation, and Test of a Multi-Model Systolic Neural-Network Accelerator</i>, Scientific Programming, Vol. 5, 1996, at 50.</p> <p>“The architecture of the implemented array and the structure of the PEs are illustrated in figure 2. Unlike the basic architecture shown in figure 1 (a), the access to the processing array is done through input and output ports located on the top-left to bottom-right diagonal. These inputs (<math>U_{in}</math> and <math>L_{in}</math>) and outputs (<math>U_{out}</math> and <math>L_{out}</math>) are present on all PEs for modularity but are used only on diagonal cells. They are always bypassed or left unconnected on non-diagonal cells. As figure 2 (b) shows, two further registers (<math>U</math> and <math>L</math>) are required for this path. The advantage is that no external hardware is required to diagonalize the inputs or un-diagonalize the outputs to and from the array (see figure 1 (a)). All the components of a vector are entered or read at the same time.</p> <p>The transposition mechanism, described by equations (15) and (16), is added to the PE by providing two control signals from the systolic instruction unit. A first signal exchanges the roles of <math>N_{in}</math> and <math>W_{in}</math>, while the other exchanges <math>S_{out}</math> and <math>E_{out}</math>, as shown in figure 3. When a computation on the transpose matrix begins, the inputs are exchanged. The quantities in the PE registers, belonging to the previous computation are output on the regular path. On the next computation step, the outputs are also exchanged. The same applies when the operation mode reverts to the direct matrix.” Paolo Ienne &amp; Marc A. Viredaz, <i>GENES IV: A Bit-Serial Processing Element for a Multi-Model Neural Network Accelerator</i>, Proceedings of the International Conference on Application-Specific Array Processors, Venice, Italy, October 1993, at 350-51.</p>

## Exhibit 1 – MANTRA I

Claim Limitation (Claim 7)	Exemplary Disclosure
	 <p>Figure 2 consists of two parts. Part (a) shows a 4x4 array of processing elements (PEs). Each PE is a square box with four input lines (labeled <math>I_1^v, I_1^h, I_2^v, I_2^h</math> for the top-left PE, and so on for the other three rows) and four output lines (<math>O_1^v, O_1^h, O_2^v, O_2^h</math> for the top-left PE, and so on). The PEs are interconnected in a mesh-like structure where each PE is connected to its four neighbors via serial lines. Part (b) shows a detailed block diagram of a single processing element. It includes an Instr. Unit, a Weight Unit (labeled D), an Arithmetic Unit, and a PS unit. The inputs are <math>Nin, WGtin, Wout, Win</math>, and the outputs are <math>Nout, INSTRout, Lin, Ein, Lout, Uout, Sin, Uin, Eout, Sout, WGtout</math>.</p> <p><b>Figure 2. (a) <math>4 \times 4</math> array architecture. (b) Processing element basic structure.</b></p> <p>Paolo Ienne &amp; Marc A. Viredaz, <i>GENES IV: A Bit-Serial Processing Element for a Multi-Model Neural Network Accelerator</i>, Proceedings of the International Conference on Application-Specific Array Processors, Venice, Italy, October 1993, at 350-51.</p> <p>“The GENES IV array is a square mesh of dedicated processing elements (PEs). Each PE holds one element of the synaptic weight matrix. It is connected by serial lines to its four neighbors as shown in figure 1(a). All input and output operations take place on the north-west to south-east diagonal. The GENES IV array implements six operations:</p> <p><b>Matrix-vector product:</b> this operation is an implementation of the classic systolic multiplier. It can alternatively be viewed as the scalar or dot product between a vector and each row of the matrix.</p> <p><b>Squared Euclidean distance:</b> using the same data flow as the matrix-vector product and a slightly modified arithmetic unit, this operation computes the squared Euclidean distance between a vector and each row of the synaptic weight matrix stored in the array.</p>

## Exhibit 1 – MANTRA I

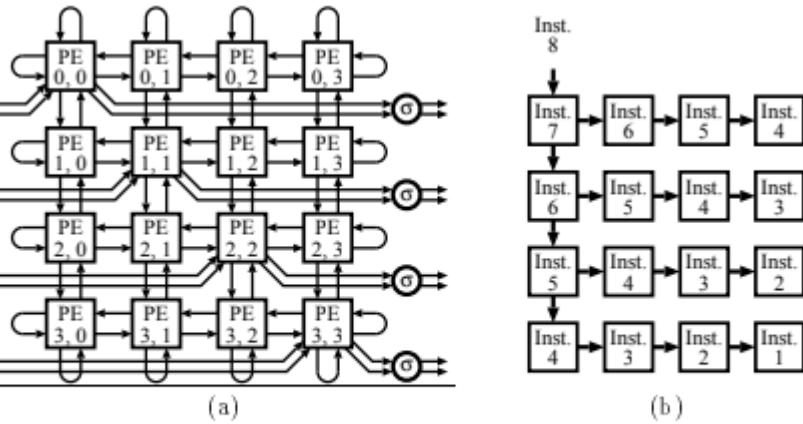
Claim Limitation (Claim 7)	Exemplary Disclosure
	<p><b>Hebbian learning rule:</b> during this operation, two vectors are injected into the array, flowing north-to-south and west-to-east respectively. The outer product of these vectors is computed, and each synaptic weight is updated with the corresponding element of the resulting matrix.</p> <p><b>Kohonen learning rule:</b> using the same data flow as the previous one, this operation updates the synaptic weights using the Kohonen learning rule.</p> <p><b>Maximum element of a vector:</b> during this operation, the same vector is injected into the array both from north to south and from west to east. At each PE, both operands are compared. The horizontal one is propagated only when it is larger than the vertical one, it is otherwise replaced by the smallest representable number <math>PS_{min}</math>. The resulting vector contains the largest element of the original vector and <math>PS_{min}</math> everywhere else.</p> <p><b>Minimum element of a vector:</b> this operation is the complement of the previous one, the smallest element of the vector being searched for.” Marc A. Viredaz &amp; Paolo Ienne, <i>MANTRA I: A Systolic Neuro-Computer</i>, In Proceedings of the International Join Conference on Neural Networks, Vol. III, Nagoya, Japan, October 1993, at 1-2.</p>  <p>Figure 1(a) shows a square systolic array of 16 Processing Elements (PEs) arranged in a 4x4 grid. Each PE is labeled with its coordinates: (0,0), (0,1), (0,2), (0,3) in the top row; (1,0), (1,1), (1,2), (1,3) in the second row; (2,0), (2,1), (2,2), (2,3) in the third row; and (3,0), (3,1), (3,2), (3,3) in the bottom row. Each PE has local connections to its neighbors and to a global output node. Figure 1(b) illustrates the instruction flow, showing a sequence of 8 instructions (Inst. 8 to Inst. 1) being processed sequentially through a series of stages, with arrows indicating the flow from one stage to the next.</p>

Figure 1: (a) Square systolic array of GENES IV PEs. (b) Instruction flow.

**Exhibit 1 – MANTRA I**

Claim Limitation (Claim 7)	Exemplary Disclosure
	<p>Marc A. Viredaz &amp; Paolo Ienne, <i>MANTRA I: A Systolic Neuro-Computer</i>, In Proceedings of the International Join Conference on Neural Networks, Vol. III, Nagoya, Japan, October 1993, at 1-2.</p> <p>“A GENES IV array is a square mesh of simple <i>processing elements (PEs)</i>. Each PE is connected by serial lines to its four neighbors as shown in figure 1. All input and output operations are performed by the PEs located on the northwest to southeast diagonal. The GENES IV structure implements six different operations grouped in three categories.” Marc A. Viredaz, <i>MANTRA I: An SIMD Processor Array for Neural Computation</i>, Spies P.P. (eds) Europäischer Informatik Kongreß Architektur von Rechensystemen Euro-ARCH, 1993, available at <a href="https://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.29.9021&amp;rep=rep1&amp;type=pdf">https://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.29.9021&amp;rep=rep1&amp;type=pdf</a>.</p>
<p>[156c] wherein the dynamic range of the possible valid inputs to the first operation is at least as wide as from 1/1,000,000 through 1,000,000 and for at least X=5% of the possible valid inputs to the first operation, the statistical mean, over repeated execution of the first operation on each specific input from the at least X% of the possible valid inputs to the first operation, of the numerical values represented by the first output signal of the LPHDR unit executing the first operation on that input differs by at least Y=0.05% from the result of an exact mathematical calculation of the first operation on the numerical values of that same input; and</p>	<p>As explained below and in the Responsive Contentions Regarding Non-Infringement and Invalidity (“Responsive Contentions”), it would have been obvious to one of skill in the art based on the disclosures in MANTRA I (alone or in combination with the teachings of Tong, Belanovic / Belanovic and Leeser, Shirazi, Atty, Lee, Sudha, Dockser, and GRAPE-3) that the dynamic range of the possible valid inputs to the first operation is at least as wide as from 1/1,000,000 through 1,000,000 and for at least X=5% of the possible valid inputs to the first operation, the statistical mean, over repeated execution of the first operation on each specific input from the at least X% of the possible valid inputs to the first operation, of the numerical values represented by the first output signal of the LPHDR unit executing the first operation on that input differs by at least Y=0.05% from the result of an exact mathematical calculation of the first operation on the numerical values of that same input.</p> <p>MANTRA I was designed to perform neural network algorithms at the minimum required precision.</p> <p>“Fixed-point arithmetic. GENES IV PEs are designed for fixed-point number representation. In the absence of general analytical results on the required precision (see also Section 4.3), simulations of the application described in Section 5 have been used to determine the precision to be implemented.” Thierry Cornu, et al. <i>Design, Implementation, and Test of a Multi-Model</i></p>

**Exhibit 1 – MANTRA I**

Claim Limitation (Claim 7)	Exemplary Disclosure
	<p><i>Systolic Neural-Network Accelerator</i>, Scientific Programming, Vol. 5, 1996, at 52.</p> <p>“Contrary to other neural networks, the Kohonen algorithm with quantized weights and inputs has received little attention so far. Three factors influencing its correct convergence can be put in evidence [17]. Clearly, there is a minimal <i>number of bits</i> required to encode the weights, depending on the input distribution and dimension, as well as the number of neurons. Second, the adaptation gain <math>\alpha</math> must decrease slowly enough, or have an initially large value, because otherwise the weight updates get rounded to zero before the algorithm has converged. Finally, the neighborhood function should decrease with the distance from the winner neuron, especially if the input dimension is low. These qualitative results were confirmed by a mathematical analysis based on the Markovian formulation of the algorithm [16], giving the necessary and sufficient conditions for the self-organization of the map in the case where the input and weight spaces are one dimensional. Roughly speaking, the results proven for the continuous case [2] also apply in the quantized case if the number of bits is large enough.” Thierry Cornu, et al. <i>Design, Implementation, and Test of a Multi-Model Systolic Neural-Network Accelerator</i>, Scientific Programming, Vol. 5, 1996, at 56.</p> <p>“It should be noticed that because of a combined effect of integer arithmetic, batch implementation, and multiple winners, MANTRA I converges with a slightly higher final error than the sequentially implemented floating-point version of the Kohonen algorithm. On a target error rate of 50% more than the minimum error of the original algorithm, the latter and the MANTRA I implementation need, respectively, 12 x 120 and 16 x 120 iterations to reach the desired error rates. This yields an algorithmic efficiency of 12/16 = 75%.</p> <p>To confirm that the algorithmic efficiency on MANTRA I is high and not very far from unity, it was tested with additional data from other applications. Test runs confirmed that the discrepancies between the MANTRA I version and the original version of the Kohonen algorithm are smaller than the standard deviation of the original algorithm itself. For instance, averaging ten runs for each of several sets of learning parameters in the speech codebook classification problem, MANTRA I actually performed better in approximately 50% of the cases.” Thierry Cornu, et al. <i>Design, Implementation, and Test of a Multi-Model Systolic Neural-Network Accelerator</i>, Scientific Programming, Vol. 5, 1996, at 58.</p>

**Exhibit 1 – MANTRA I**

Claim Limitation (Claim 7)	Exemplary Disclosure
	<p>“In conclusion, the final rate reached by MANTRA I appears acceptable for the power-system application and the number of iterations required is approximately equivalent to that needed by a traditional floating-point version of the Kohonen algorithm.” Thierry Cornu, et al. <i>Design, Implementation, and Test of a Multi-Model Systolic Neural-Network Accelerator</i>, Scientific Programming, Vol. 5, 1996, at 59.</p> <p>“The choice of the arithmetic precision of the hardware is quite delicate. A first, typical choice, motivated by hardware simplicity, is to represent all variables in two’s complement fixed point. As for the required number of bits for each element, theoretical analysis usually leads to absolute minimum boundaries that, on themselves, cannot guarantee convergence [13]. In the absence of strong analytical grounds, simulation usually provides the designer with the required information. In [2] and [7] simulations are performed to determine the needs of typical backpropagation applications. These have been completed by simulations of the Kohonen model in the application that prompted the development of the present system.” Paolo Ienne &amp; Marc A. Viredaz, <i>GENES IV: A Bit-Serial Processing Element for a Multi-Model Neural Network Accelerator</i>, Proceedings of the International Conference on Application-Specific Array Processors, Venice, Italy, October 1993, at 352.</p> <p>Reduced precision computations were commonly used in connection with neurocomputers like the MANTRA I. <i>See, e.g.:</i></p> <p>“Another simplification made possible by ANN algorithms depends on the reduced precision required in most calculation and in most models (Asanović and Morgan 1991; Holt and Baker 1991; Holt and Hwang 1993; Thiran et al. 1994). As a consequence, the designer may avoid area-expensive floating-point arithmetic units and use reduced integer precision (all but one system among those cited in the following sections use fixed-point arithmetic). This results in arithmetic units, registers and data paths using less area on the die and in a reduction of the physical resources devoted to communication (e.g., less IC pins).” Paolo Ienne Lopez, <i>Programmable VLSI Systolic Processors for Neural Network and Matrix Computations</i> (unpublished Ph. D. dissertation, École Polytechnique Fédérale de Lausanne), 1996, at 9-10.</p>

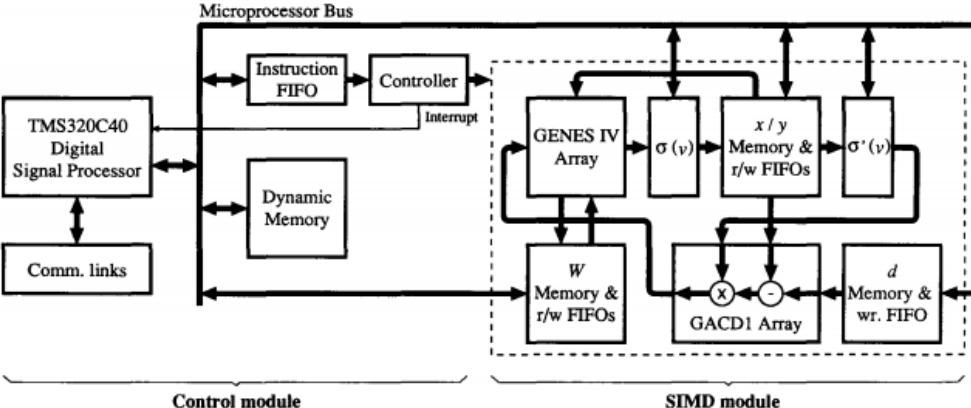
**Exhibit 1 – MANTRA I**

Claim Limitation (Claim 7)	Exemplary Disclosure
	<p>“Other systems try to take advantage of other peculiarities of ANN algorithms, such as a reduced precision required in the computations. This makes it possible to develop ad-hoc processing elements which are characterized by a small size and cost. In turn, this often enables one to design systems with a very high processing element count and therefore to increase the degree of parallelism.” Paolo Ienne, <i>Digital Systems for Neural Networks</i>, PROC. SPIE 10279, Digital Signal Processing Technology: A Critical Review, April 25, 1995, at 11.</p> <p>Accordingly, MANTRA I discloses the use of arithmetic units designed to perform calculations using reduced-precision 16-bit fixed-point math for calculations that typically use 32-bit floating point math, operating on only the 16 most significant bits in the registers. <i>See, e.g.:</i></p> <p>“The registers <math>W_0</math> and <math>W_1</math> are 33 bits wide (including an overflow bit), but only the most significant 16 bits are used for non-learning operations. The other registers D, PS, U, and L are all 40 bits wide. Only 16 bits of the register D are however used for other operations than the search for the largest/smallest element of a vector. Similarly, the register PS is used as a 17-bit value for the two learning operations.” Marc A. Viredaz, <i>MANTRA I: An SIMD Processor Array for Neural Computation</i>, Spies P.P. (eds) Europäischer Informatik Kongreß Architektur von Rechensystemen Euro-ARCH, 1993, available at <a href="https://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.29.9021&amp;rep=rep1&amp;type=pdf">https://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.29.9021&amp;rep=rep1&amp;type=pdf</a>.</p> <p>“The computation is performed on signed fixed-point values. The inputs and the weights are coded on 16 bits. The weights have 16 additional bits. but these are used only during learning (weight update operations). Outputs are computed on 40 bits.” Thierry Cornu, et al. <i>Design, Implementation, and Test of a Multi-Model Systolic Neural-Network Accelerator</i>, Scientific Programming, Vol. 5, 1996, at 50.</p> <p>“BP implementations typically use 32-bit floating point math. This largely eliminates scaling, precision and dynamic range issues. Efficient hardware implementation dictates integer arithmetic units with precision no greater than required. Baker [Bak90] has shown 16-bit integer weights are sufficient for BP training and much lower values adequate for use after training.” Hal McCartor, <i>Back Propagation Implementation on the Adaptive Solutions CNAPS</i></p>

**Exhibit 1 – MANTRA I**

Claim Limitation (Claim 7)	Exemplary Disclosure
	<p><i>Neurocomputer Chip</i>, ADVANCES IN NEURAL INFORMATION PROCESSING SYSTEMS 3, 1990, at 1029.</p> <p>For the reasons explained in the Responsive Contentions, it would have been obvious to one of skill in the art to have substituted the fixed-point number format used in MANTRA I for a floating-point format that met the claimed minimum range and precision requirements, and to have used the reduced-precision floating-point number formats disclosed in Tong, Dockser, Belanovic / Belanovic and Leeser, Lee, Shirazi, Sudha, Aty, and TMS 320C32 or the logarithmic format disclosed in GRAPE-3 and Hoefflinger, either alone or in combination. <i>See also</i> Appendix to Responsive Contentions (detailing error rates associated with different mantissa sizes).</p>
<p>[156d] at least one first computing device adapted to control the operation of the at least one first LPHDR execution unit;</p>	<p>MANTRA I discloses at least one first computing device adapted to control the operation of the at least one first LPHDR execution unit. <i>See, e.g.:</i></p> <p>“The MANTRA I machine is controlled by a TMS320C40 digital signal processor (DSP) from Texas Instruments. Two of its six eight-bit built-in communication links connect the machine to another TMS320C40 processor inside a SUN SPARCstation (Fig. 2). From a software point of view, the intermediate DSP is transparent. The MANTRA I machine (the systolic array and its control processor) is completely controlled by the front-end workstation but could be easily integrated into any other computer system based on TMS320C40 processors.” Thierry Cornu, et al. <i>Design, Implementation, and Test of a Multi-Model Systolic Neural-Network Accelerator</i>, Scientific Programming, Vol. 5, 1996, at 48.</p> <p>“The structure of the MANTRA I system [18] is shown in Figure 3. The <i>control module</i> is the SISD system based on the DSP. It controls the <i>parallel</i> or <i>SIMD module</i> by dispatching horizontally coded instructions through an FIFO. The SIMD module is frozen when no instruction is pending.” Thierry Cornu, et al. <i>Design, Implementation, and Test of a Multi-Model Systolic Neural-Network Accelerator</i>, Scientific Programming, Vol. 5, 1996, at 49.</p>

## Exhibit 1 – MANTRA I

Claim Limitation (Claim 7)	Exemplary Disclosure
	 <p>FIGURE 3 Architecture of the MANTRA I machine.</p> <p>Thierry Cornu, et al. <i>Design, Implementation, and Test of a Multi-Model Systolic Neural-Network Accelerator</i>, Scientific Programming, Vol. 5, 1996, at 48.</p> <p>“The MANTRA I system [14] is shown in figure 6. The computational heart is a GENES IV array of 40 x 40 PEs. The sequencing of the systolic array is performed by a TMS320C40 digital signal processor (DSP) from Texas Instruments. It also handles the communications with a host workstation and between different MANTRA I machines connected through the dedicated communication links of the DSP.” Paolo Ienne &amp; Marc A. Viredaz, <i>GENES IV: A Bit-Serial Processing Element for a Multi-Model Neural Network Accelerator</i>, Proceedings of the International Conference on Application-Specific Array Processors, Venice, Italy, October 1993, at 354.</p> <p>“The control part is a complete SISD system based on a commercial microprocessor: the TMS320C40 from Texas Instruments. It configures the SIMD part, dispatches instructions, and manages the inputs and outputs. It also handles the communications with a host computer and between different interconnected MANTRA I computers.” Marc A. Viredaz &amp; Paolo Ienne, <i>MANTRA I: A Systolic Neuro-Computer</i>, In Proceedings of the International Join Conference on Neural Networks, Vol. III, Nagoya, Japan, October 1993, at 3.</p>

**Exhibit 1 – MANTRA I**

Claim Limitation (Claim 7)	Exemplary Disclosure
	<p>“The control part of the MANTRA I machine, groups all the units shown outside the dashed box of figure 4. Its tasks are to configure the SIMD part, to dispatch instructions, and to manage the inputs and outputs. This part is a complete SISD system based on the TMS320C40 microprocessor from Texas Instruments running at 20 MHz. It should also handle the communications with a host computer and between different interconnected MANTRA I computers.” Marc A. Viredaz, <i>MANTRA I: An SIMD Processor Array for Neural Computation</i>, Spies P.P. (eds) Europäischer Informatik Kongreß Architektur von Rechensystemen Euro-ARCH, 1993, available at <a href="https://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.29.9021&amp;rep=rep1&amp;type=pdf">https://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.29.9021&amp;rep=rep1&amp;type=pdf</a>.</p>

## Exhibit 1 – MANTRA I

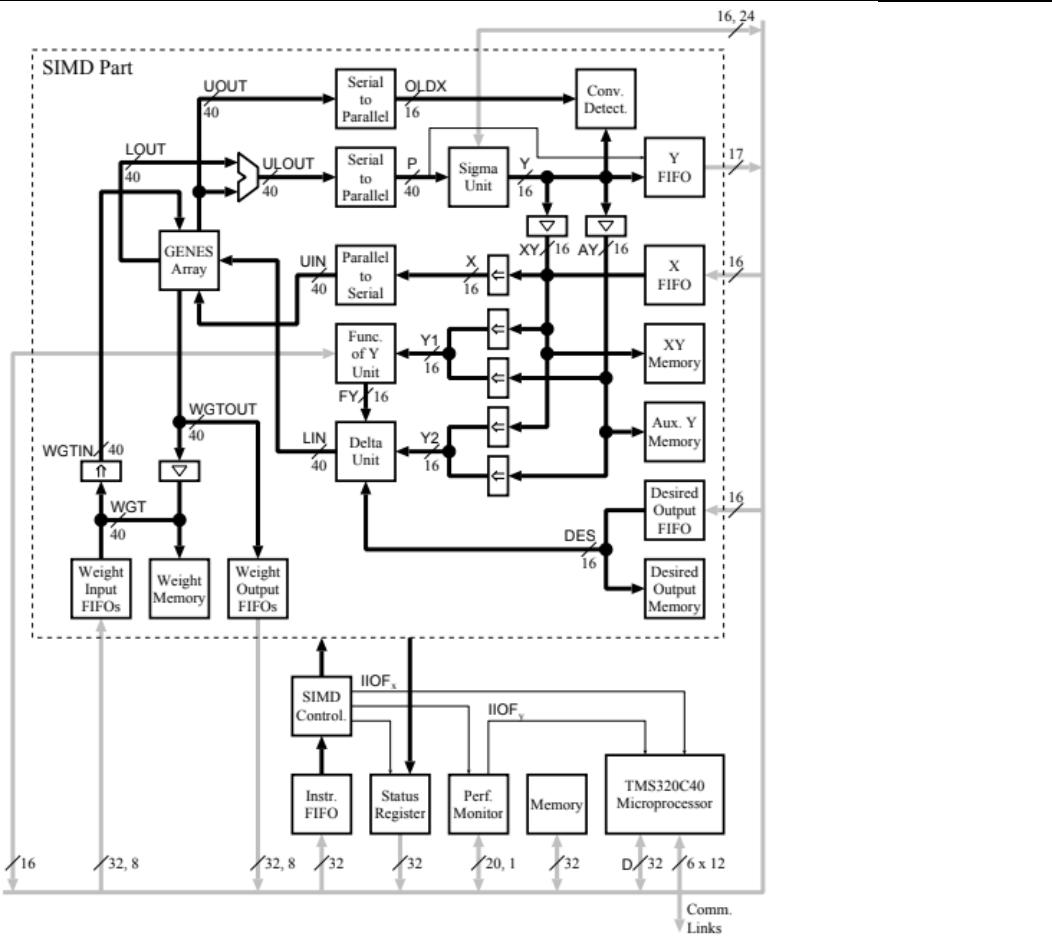
Claim Limitation (Claim 7)	Exemplary Disclosure
	 <p>The diagram illustrates the architecture of the MANTRA I machine. It is organized into two main sections: the SIMD Part and the Exemplary Disclosure section.</p> <p><b>SIMD Part:</b> This section contains several key components:     <ul style="list-style-type: none"> <li><b>GENES Array:</b> A central component that receives input from Weight Input FIFOs and Weight Memory, and provides output to Weight Output FIFOs.</li> <li><b>Serial to Parallel converters:</b> Convert data between parallel and serial formats.</li> <li><b>Sigma Unit:</b> A unit that processes data from the GENES Array and the Delta Unit.</li> <li><b>Conv. Detect. unit:</b> Detects convolutional features.</li> <li><b>FIFOs and Memories:</b> Various FIFOs (X, Y, XY, Aux. Y) and memories (Desired Output FIFO/Memory) for temporary storage and data exchange.</li> </ul> </p> <p><b>Exemplary Disclosure:</b> This section shows the overall system architecture:     <ul style="list-style-type: none"> <li><b>SIMD Control:</b> Manages the SIMD operations.</li> <li><b>Instr. FIFO:</b> Stores instructions.</li> <li><b>Status Register:</b> Monitors system status.</li> <li><b>Perf. Monitor:</b> Measures performance metrics.</li> <li><b>Memory:</b> Provides memory access.</li> <li><b>TMS320C40 Microprocessor:</b> Handles general control.</li> </ul>     The system is interconnected via various buses and communication links, with data paths labeled with bit widths such as 16, 24, 32, 40, and 6 x 12 bits. </p>

Figure 4: Architecture of the MANTRA I machine.

Marc A. Viredaz, *MANTRA I: An SIMD Processor Array for Neural Computation*, Spies P.P. (eds) Europäischer Informatik Kongreß Architektur von Rechensystemen Euro-ARCH, 1993, available at

<https://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.29.9021&rep=rep1&type=pdf>

**Exhibit 1 – MANTRA I**

Claim Limitation (Claim 7)	Exemplary Disclosure
<p>[156e] wherein the at least one first computing device comprises at least one of a central processing unit (CPU), a graphics processing unit (GPU), a field programmable gate array (FPGA), a microcode-based processor, a hardware sequencer, and a state machine;</p>	<p>MANTRA I discloses at least one first computing device comprises at least one of a central processing unit (CPU), a graphics processing unit (GPU), a field programmable gate array (FPGA), a microcode-based processor, a hardware sequencer, and a state machine. <i>See, e.g.:</i></p> <p>“The MANTRA I machine is controlled by a TMS320C40 digital signal processor (DSP) from Texas Instruments. Two of its six eight-bit built-in communication links connect the machine to another TMS320C40 processor inside a SUN SPARCstation (Fig. 2). From a software point of view, the intermediate DSP is transparent. The MANTRA I machine (the systolic array and its control processor) is completely controlled by the front-end workstation but could be easily integrated into any other computer system based on TMS320C40 processors.” Thierry Cornu, et al. <i>Design, Implementation, and Test of a Multi-Model Systolic Neural-Network Accelerator</i>, Scientific Programming, Vol. 5, 1996, at 48.</p> <p>“The MANTRA I system [14] is shown in figure 6. The computational heart is a GENES IV array of 40 x 40 PEs. The sequencing of the systolic array is performed by a TMS320C40 digital signal processor (DSP) from Texas Instruments. It also handles the communications with a host workstation and between different MANTRA I machines connected through the dedicated communication links of the DSP.” Paolo Ienne &amp; Marc A. Viredaz, <i>GENES IV: A Bit-Serial Processing Element for a Multi-Model Neural Network Accelerator</i>, Proceedings of the International Conference on Application-Specific Array Processors, Venice, Italy, October 1993, at 354.</p> <p>“The control part is a complete SISD system based on a commercial microprocessor: the TMS320C40 from Texas Instruments. It configures the SIMD part, dispatches instructions, and manages the inputs and outputs. It also handles the communications with a host computer and between different interconnected MANTRA I computers.” Marc A. Viredaz &amp; Paolo Ienne, <i>MANTRA I: A Systolic Neuro-Computer</i>, In Proceedings of the International Join Conference on Neural Networks, Vol. III, Nagoya, Japan, October 1993, at 3.</p> <p>“The control part of the MANTRA I machine, groups all the units shown outside the dashed box of figure 4. Its tasks are to configure the SIMD part, to dispatch instructions, and to manage the inputs and outputs. This part is a complete SISD system based on the TMS320C40</p>

**Exhibit 1 – MANTRA I**

Claim Limitation (Claim 7)	Exemplary Disclosure
	<p>microprocessor from Texas Instruments running at 20 MHz. It should also handle the communications with a host computer and between different interconnected MANTRA I computers.” Marc A. Viredaz, <i>MANTRA I: An SIMD Processor Array for Neural Computation</i>, Spies P.P. (eds) Europäischer Informatik Kongreß Architektur von Rechensystemen Euro-ARCH, 1993, available at <a href="https://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.29.9021&amp;rep=rep1&amp;type=pdf">https://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.29.9021&amp;rep=rep1&amp;type=pdf</a>.</p>

## Exhibit 1 – MANTRA I

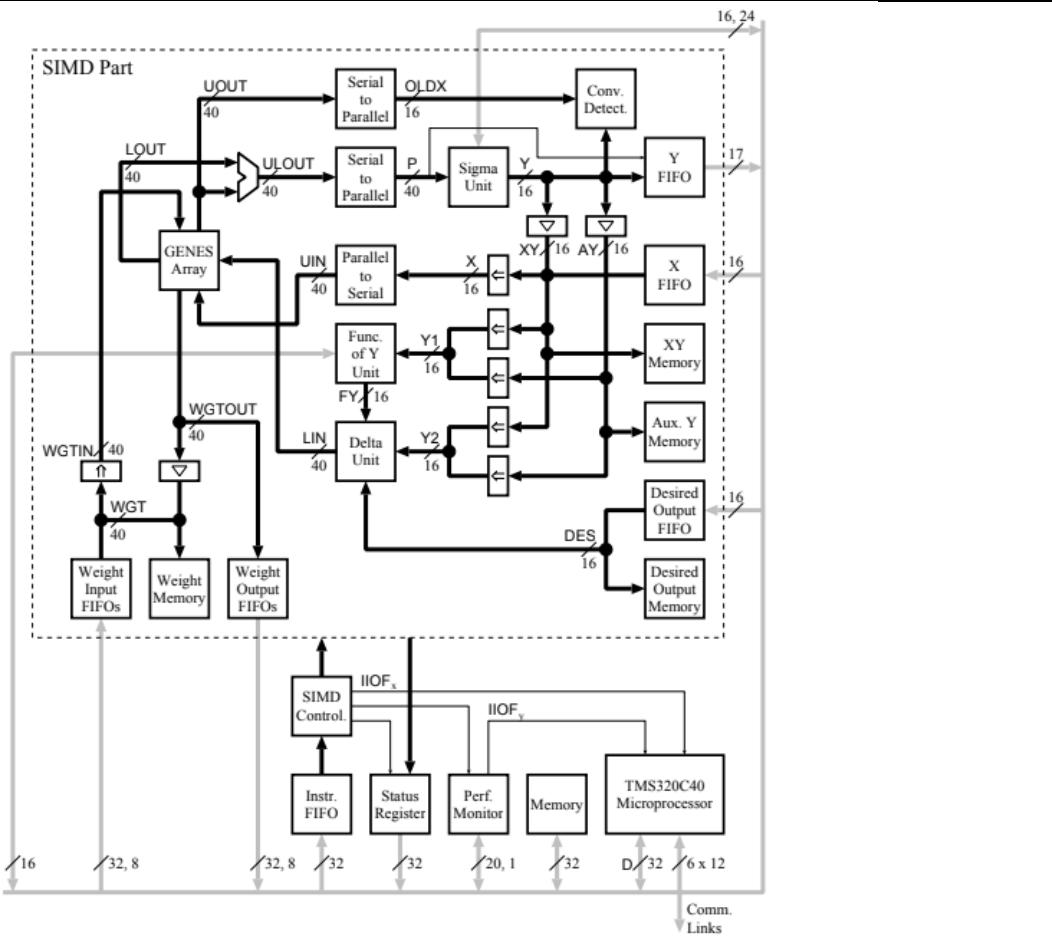
Claim Limitation (Claim 7)	Exemplary Disclosure
	 <p>The diagram illustrates the architecture of the MANTRA I machine. It is organized into two main sections: the SIMD Part and the Exemplary Disclosure section.</p> <p><b>SIMD Part:</b> This section contains several functional blocks:     <ul style="list-style-type: none"> <li><b>GENES Array:</b> Receives inputs from <b>WGTIN</b> (40 bits) and <b>UOUT</b> (40 bits).</li> <li><b>Serial to Parallel converters:</b> Convert data from serial to parallel formats.</li> <li><b>Sigma Unit:</b> Processes data from the GENES Array and the Serial to Parallel converter.</li> <li><b>Conv. Detect. unit:</b> Detects convolutional features.</li> <li><b>FIFOs and Memories:</b> <b>X FIFO</b>, <b>XY Memory</b>, <b>Aux. Y Memory</b>, <b>Desired Output FIFO</b>, and <b>Desired Output Memory</b>.</li> </ul> </p> <p><b>Exemplary Disclosure:</b> This section shows the overall system architecture:     <ul style="list-style-type: none"> <li><b>SIMD Control:</b> Manages the SIMD operations.</li> <li><b>Instr. FIFO:</b> Stores instructions.</li> <li><b>Status Register:</b> Monitors system status.</li> <li><b>Perf. Monitor:</b> Measures performance metrics.</li> <li><b>Memory:</b> Provides memory access.</li> <li><b>TMS320C40 Microprocessor:</b> Handles general control.</li> <li><b>Communication Links:</b> Connects the SIMD Control to the SIMD Part and the Exemplary Disclosure section.</li> </ul> </p> <p>Arrows indicate data flow between components, with bit widths labeled on many of the connections. For example, <b>WGTIN</b> is 40 bits, <b>UOUT</b> is 40 bits, and <b>Serial to Parallel</b> conversions are 40 bits wide. The <b>Desired Output</b> path is 16 bits wide.</p>

Figure 4: Architecture of the MANTRA I machine.

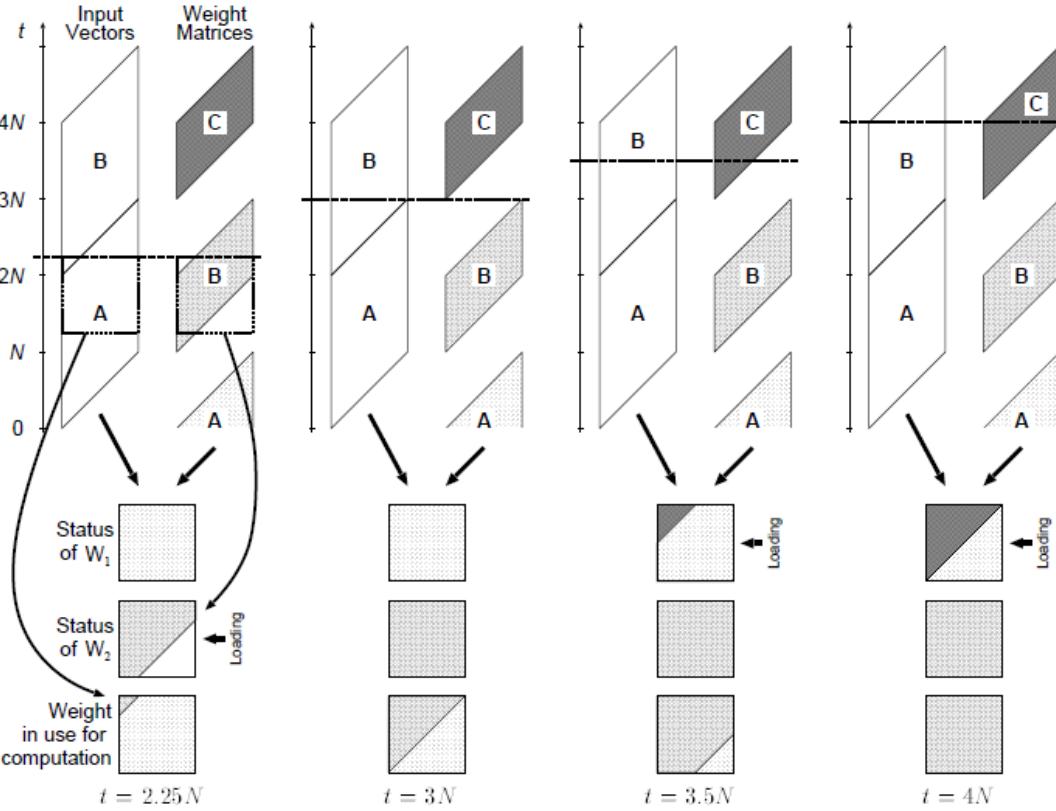
Marc A. Viredaz, *MANTRA I: An SIMD Processor Array for Neural Computation*, Spies P.P. (eds) Europäischer Informatik Kongreß Architektur von Rechensystemen Euro-ARCH, 1993, available at

<https://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.29.9021&rep=rep1&type=pdf>

**Exhibit 1 – MANTRA I**

Claim Limitation (Claim 7)	Exemplary Disclosure
<p>[156f] and, wherein the number of LPHDR execution units in the device exceeds by at least one hundred the non-negative integer number of execution units in the device adapted to execute at least the operation of multiplication on floating point numbers that are at least 32 bits wide.</p>	<p>MANTRA I discloses the number of LPHDR execution units in the device exceeds by at least one hundred the non-negative integer number of execution units in the device adapted to execute at least the operation of multiplication on floating point numbers that are at least 32 bits wide.</p> <p>MANTRA I incorporates 400 to 1,600 processing elements. <i>See, e.g.:</i></p> <p>“The machine, which is currently under test, consists of four printed circuit boards : (1) the processor board (34 chips), (2) the control board (135 chips), (3) the input/output board (465 chips including 14 GACD1 chips), (4) the GENES IV array board (containing a matrix of 10 x 10 = 100 GENES IV chips). One of the latter board is required for configurations of the machine up to 20 x 20 = 400 PEs. For larger configurations—up to 40 x 40 = 1600 PEs—four such boards should be interconnected. A more detailed description of the MANTRA I machine can be found in [vir93].” Marc A. Viredaz &amp; Paolo Ienne, <i>MANTRA I: A Systolic Neuro-Computer</i>, In Proceedings of the International Join Conference on Neural Networks, Vol. III, Nagoya, Japan, October 1993, at 3.</p> <p>“The heart of the SIMD part in the MANTRA I machine is an array of 40 x 40 GENES IV PEs running at 10 MHz.” Marc A. Viredaz &amp; Paolo Ienne, <i>MANTRA I: A Systolic Neuro-Computer</i>, In Proceedings of the International Join Conference on Neural Networks, Vol. III, Nagoya, Japan, October 1993, at 3.</p>

## Exhibit 1 – MANTRA I

Claim Limitation (Claim 7)	Exemplary Disclosure
	 <p>Figure 3.11: Genes IV background weight exchange mechanism. The letters <b>A</b>, <b>B</b>, and <b>C</b> identify each weight matrix and the corresponding operand vectors.</p> <p>Paolo Ienne Lopez, <i>Programmable VLSI Systolic Processors for Neural Network and Matrix Computations</i> (unpublished Ph. D. dissertation, École Polytechnique Fédérale de Lausanne), 1996, at 44.</p> <p>“The dashed box delimits the parallel or SIMD part of the machine, whose computational heart is an array of up to 40 x 40 Genes IV PEs. The control part is a complete <i>single instruction-single data-stream</i> (SISD) system based on the TMS320C40 DSP from TEXAS INSTRUMENTS, with its own memory. This processor has six built-in communication channels,</p>

**Exhibit 1 – MANTRA I**

Claim Limitation (Claim 7)	Exemplary Disclosure
	two of which have been used in the prototype to connect Mantra I to a SUN MICROSYSTEMS workstation where data and programs are stored. The workstation also provides the user interface. Cornu et al. (1996) have described the software upper layer that implements the user front-end.” Paolo Ienne Lopez, <i>Programmable VLSI Systolic Processors for Neural Network and Matrix Computations</i> (unpublished Ph. D. dissertation, École Polytechnique Fédérale de Lausanne), 1996, at 44.

**Exhibit 1 – MANTRA I**

'273 Patent

<b>Claim Limitation (Claim 53)</b>	<b>Exemplary Disclosure</b>
[273a] A device:	MANTRA I discloses a device. Specifically, the MANTRA I is a neurocomputer. <i>See [156a].</i>
[273b] comprising at least one first low precision high-dynamic range (LPHDR) execution unit adapted to execute a first operation on a first input signal representing a first numerical value to produce a first output signal representing a second numerical value,	MANTRA I discloses at least one first low precision high dynamic range (LPHDR) execution unit adapted to execute a first input signal representing a first numerical value to produce a first output signal representing a second numerical value. <i>See [156b]</i>
[273c] wherein the dynamic range of the possible valid inputs to the first operation is at least as wide as from 1/1,000,000 through 1,000,000 and for at least X=5% of the possible valid inputs to the first operation, the statistical mean, over repeated execution of the first operation on each specific input from the at least X% of the possible valid inputs to the first operation, of the numerical values represented by the first output signal of the LPHDR unit executing the first operation on that input differs by at least Y=0.05% from the result of an exact mathematical calculation of the first operation on the numerical values of that same input. <i>See [156c]; see also Appendix to Responsive Contentions (detailing error rates associated with different mantissa sizes).</i>	MANTRA I discloses the dynamic range of the possible valid inputs to the first operation is at least as wide as from 1/1,000,000 through 1,000,000 and for at least X=5% of the possible valid inputs to the first operation, the statistical mean, over repeated execution of the first operation on each specific input from the at least X% of the possible valid inputs to the first operation, of the numerical values represented by the first output signal of the LPHDR unit executing the first operation on that input differs by at least Y=0.05% from the result of an exact mathematical calculation of the first operation on the numerical values of that same input. <i>See [156c]; see also Appendix to Responsive Contentions (detailing error rates associated with different mantissa sizes).</i>

**Exhibit 1 – MANTRA I**

<b>Claim Limitation (Claim 53)</b>	<b>Exemplary Disclosure</b>
[273d] wherein the number of LPHDR execution units in the device exceeds by at least one hundred the non-negative integer number of execution units in the device adapted to execute at least the operation of multiplication on floating point numbers that are at least 32 bits wide.	MANTRA I discloses the number of LPHDR execution units in the device exceeds by at least one hundred the non-negative integer number of execution units in the device adapted to execute at least the operation of multiplication on floating point numbers that are at least 32 bits wide. <i>See [156f].</i>

**Exhibit 1 – MANTRA I**

'961 Patent

<b>Claim Limitation (Claim 4)</b>	<b>Exemplary Disclosure</b>
[961a] A device comprising:	MANTRA I discloses a device. Specifically, the MANTRA I is a neurocomputer. <i>See [156a].</i>
[961b] at least one first low precision high-dynamic range (LPHDR) execution unit adapted to execute a first operation on a first input signal representing a first numerical value to produce a first output signal representing a second numerical value,	MANTRA I discloses at least one first low precision high dynamic range (LPHDR) execution unit adapted to execute a first input signal representing a first numerical value to produce a first output signal representing a second numerical value. <i>See [156b].</i>
[961c] wherein the dynamic range of the possible valid inputs to the first operation is at least as wide as from 1/1,000,000 through 1,000,000 and for at least X=10% of the possible valid inputs to the first operation, the statistical mean, over repeated execution of the first operation on each specific input from the at least X% of the possible valid inputs to the first operation, of the numerical values represented by the first output signal of the LPHDR unit executing the first operation on that input differs by at least Y=0.2% from the result of an exact mathematical calculation of the first operation on the numerical values of that same input. <i>See [156c]; see also Appendix to Responsive Contentions (detailing error rates associated with different mantissa sizes).</i>	MANTRA I discloses the dynamic range of the possible valid inputs to the first operation is at least as wide as from 1/1,000,000 through 1,000,000 and for at least X=10% of the possible valid inputs to the first operation, the statistical mean, over repeated execution of the first operation on each specific input from the at least X% of the possible valid inputs to the first operation, of the numerical values represented by the first output signal of the LPHDR unit executing the first operation on that input differs by at least Y=0.2% from the result of an exact mathematical calculation of the first operation on the numerical values of that same input. <i>See [156c]; see also Appendix to Responsive Contentions (detailing error rates associated with different mantissa sizes).</i>
[961d] at least one first computing device adapted to control the	MANTRA I discloses at least one first computing device adapted to control the operation of the at least one first LPHDR execution unit. <i>See [156d].</i>

**Exhibit 1 – MANTRA I**

<b>Claim Limitation (Claim 4)</b>	<b>Exemplary Disclosure</b>
operation of the at least one first LPHDR execution unit.	
<b>Claim Limitation (Claim 13)</b>	<b>Exemplary Disclosure</b>
[961e] A device comprising:	MANTRA I discloses a device. Specifically, the MANTRA I is a neurocomputer. <i>See [156a].</i>
[961f] a plurality of components comprising:	<i>See</i> MANTRA I discloses at least one first low precision high dynamic range (LPHDR) execution unit adapted to execute a first input signal representing a first numerical value to produce a first output signal representing a second numerical value. <i>See [156b]. See also</i> MANTRA I discloses at least one first computing device adapted to control the operation of the at least one first LPHDR execution unit. <i>See above [156d].</i>
[961g] at least one first low precision high-dynamic range (LPHDR) execution unit adapted to execute a first operation on a first input signal representing a first numerical value to produce a first output signal representing a second numerical value,	MANTRA I discloses at least one first low precision high dynamic range (LPHDR) execution unit adapted to execute a first input signal representing a first numerical value to produce a first output signal representing a second numerical value. <i>See [156b].</i>
[961h] wherein the dynamic range of the possible valid inputs to the first operation is at least as wide as from 1/1,000,000 through 1,000,000 and for at least X=10% of the possible valid inputs to the first operation, the statistical mean, over repeated execution of the first operation on each specific input from the at least X% of the possible valid inputs to the first operation, of the numerical values represented by the first output signal of the LPHDR unit executing the first operation on that input differs by at least Y=0.2% from the result of an exact mathematical calculation of the first operation on the numerical values of that same input. <i>See [156c]; see also</i> Appendix to Responsive Contentions (detailing error rates associated with different mantissa sizes).	MANTRA I discloses the dynamic range of the possible valid inputs to the first operation is at least as wide as from 1/1,000,000 through 1,000,000 and for at least X=10% of the possible valid inputs to the first operation, the statistical mean, over repeated execution of the first operation on each specific input from the at least X% of the possible valid inputs to the first operation, of the numerical values represented by the first output signal of the LPHDR unit executing the first operation on that input differs by at least Y=0.2% from the result of an exact mathematical calculation of the first operation on the numerical values of that same input. <i>See [156c]; see also</i> Appendix to Responsive Contentions (detailing error rates associated with different mantissa sizes).

**Exhibit 1 – MANTRA I**

<b>Claim Limitation (Claim 13)</b>	<b>Exemplary Disclosure</b>
operation on that input differs by at least Y=0.2% from the result of an exact mathematical calculation of the first operation on the numerical values of that same input.	